

The functional mean-shift algorithm for mode hunting and clustering in infinite dimensions

Mattia Ciollaro
 Christopher R. Genovese
 Jing Lei
 Larry Wasserman

*Department of Statistics
 Carnegie Mellon University
 5000 Forbes Avenue
 Pittsburgh (PA)
 15213 USA*

CIOLLARO@CMU.EDU
 GENOVESE@STAT.CMU.EDU
 JINGLEI@ANDREW.CMU.EDU
 LARRY@STAT.CMU.EDU

Editor: ADD EDITOR

Abstract

We introduce the *functional mean-shift algorithm*, an iterative algorithm for estimating the local modes of a surrogate density from functional data. We show that the algorithm can be used for cluster analysis of functional data. We propose a test based on the bootstrap for the significance of the estimated local modes of the surrogate density. We present two applications of our methodology. In the first application, we demonstrate how the functional mean-shift algorithm can be used to perform *spike sorting*, i.e. cluster neural activity curves. In the second application, we use the functional mean-shift algorithm to distinguish between original and fake signatures.

Keywords: Clustering, Functional data analysis, mean-shift algorithm, Mode hunting, Surrogate density

1. Introduction

A probability density function contains information that can be used for clustering data drawn from that distribution. Figure 1 depicts a probability density function that is a mixture of two distinct unimodal densities. Intuitively, we can say that this density represents two clusters. To unambiguously characterize the two clusters, we can use the *basin of attraction* of each local mode: if we repeatedly shift a point x_0 on the x axis along the direction of steepest ascent, the sequence thus generated converges either to the left or to the right local mode. If the destination of x_0 is the left (right) local mode, then x_0 belongs to the basin of attraction of the left (right) local mode. The set of points in the domain of the density whose common destination is the left (right) local mode naturally form a well-defined cluster. This idea of *modal clustering* easily generalizes to q -dimensional densities (Figure 2). If the density is unknown, as is usually the case in statistical inference, a set of empirical local modes and the corresponding set of empirical clusters can be obtained by performing the above procedure on an estimate of the density.




Figure 1: A probability density function representing two clusters. Each cluster is the basin of attraction of a local mode.

In this paper, we extend the idea of modal clustering to non-Euclidean spaces of infinite dimension. We demonstrate that:

1. We can meaningfully define a surrogate notion of density in an infinite dimensional space that lacks a natural dominating measure.
2. We can find the local modes of this surrogate density and the corresponding clusters.
3. We test if, given a sample of infinite dimensional data, a set of estimated local modes and their corresponding clusters are significant or only due to random fluctuation.

In particular,

1. We consider a surrogate of the notion of probability density that does not require the existence of a dominating measure. Suppose that the probability associated to the open ball of radius ϵ centered at the element x of the infinite dimensional space \mathcal{X} admits the asymptotic approximation

$$P(X \in B(x, \epsilon)) = p(x)\phi(\epsilon) + o(\phi(\epsilon))$$

for a functional p and a function $\phi(\epsilon)$ as $\epsilon \rightarrow 0$. Then, we call the functional p a surrogate density for P . This p can be consistently estimated.

2. We generalize the *mean-shift algorithm*, a simple and fast iterative procedure that estimates the local modes of a probability density, to functional data and surrogate densities. The functional mean-shift algorithm repeatedly shift a point $x \in \mathcal{X}$ towards its closest local sample mean using the update equation

$$x \leftarrow \frac{\sum_{X \in S} k\left(\frac{d(X, x)}{h}\right) X}{\sum_{X \in S} k\left(\frac{d(X, x)}{h}\right)},$$

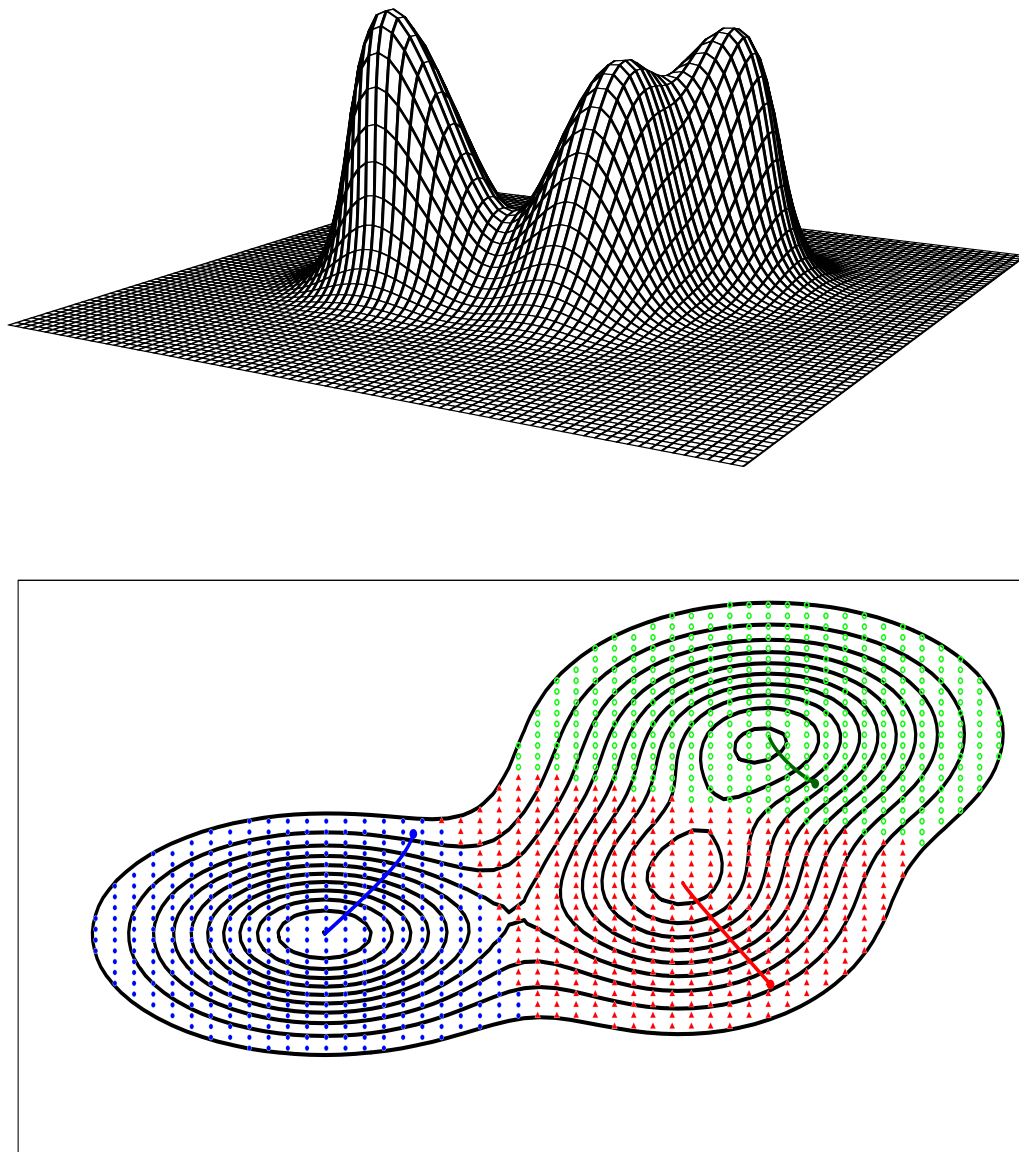


Figure 2: Modal clustering on a bivariate density with compact support. The density has three local modes and therefore three well-defined clusters. Within each basin of attraction, the gradient ascent path of a point that belongs to that basin is highlighted.

where X is an element of the sample $\mathcal{S} = \{X_1, \dots, X_n\} \subset \mathcal{X}$, k is a kernel function, d is a suitable distance, and $h > 0$ is a bandwidth parameter. The sequence generated by the above update equation converges to a critical point x^* of a kernel density estimate of the unknown surrogate density p .

3. We construct a bootstrap test for the significance of the estimated local modes. Suppose that x^* is a candidate local mode of the unknown surrogate density p . We test whether the curvature of the surrogate density p at x^* is negative:

$$H_0 : \lambda_{x^*} \geq 0$$

vs.

$$H_1 : \lambda_{x^*} < 0,$$

where $\lambda_{x^*} = \sup_{\|y\|=1} p_{x^*}^{(2)}(y, y)$ and $p_{x^*}^{(2)}$ is the Hessian of the surrogate density p at x^* .

We illustrate our methodology and its performance with two applications on smooth curves. The first application aims at clustering a set of curves corresponding to the neural activity of a Macaque monkey performing a particular task. In the second application, we apply our methodology to distinguish between original and fake signatures.

2. Related literature

The problem of estimating the mode of an unknown probability density dates back at least as early as [Parzen \(1962\)](#). The work of Parzen has been extended in several directions, for instance by [Chernoff \(1964\)](#), who shows the asymptotic normality of the sample mode, [Eddy \(1980\)](#), who studies its rate of convergence, and [Romano \(1988\)](#), who derives asymptotic minimax risk lower bounds. [Donoho and Liu \(1991\)](#) obtain the minimax rate for estimating a mode in one dimension. [Vieu \(1996\)](#) proposes four density mode estimators that attain the same rate of convergence of [Eddy \(1980\)](#). [Klemelä \(2005\)](#) proposes an adaptive estimator which attains the optimal rate of convergence.

The more general problem of estimating nonparametrically the local modes of an unknown probability density has its roots in the seminal work of [Fukunaga and Hostetler \(1975\)](#) on the *mean-shift algorithm* (see also [Silverman, 1981](#)). The work of [Fukunaga and Hostetler \(1975\)](#), later generalized by [Cheng \(1995\)](#), also promotes the mean-shift algorithm as a tool for cluster analysis ([Silverman, 1986](#)). More recently, the algorithm has been widely used in the computer science community for image segmentation tasks ([Comaniciu and Meer, 2002](#) and [Carreira-Perpiñán, 2006](#)).

The mechanics of the mean-shift algorithm are simple. An arbitrary point $x \in \mathbb{R}^q$ is repeatedly shifted towards its closest local sample mean by the iterated application of the mean-shift update equation

$$x \leftarrow \frac{\sum_{X \in \mathcal{S}} k\left(\frac{\|X-x\|}{h}\right) X}{\sum_{X \in \mathcal{S}} k\left(\frac{\|X-x\|}{h}\right)}, \quad (1)$$

where X is a point belonging to the sample $\mathcal{S} = \{X_1, \dots, X_n\} \subset \mathbb{R}^q$, k is a kernel function, $\|\cdot\|$ denotes the usual Euclidean norm, and $h > 0$ is a bandwidth parameter. The update is

performed iteratively until convergence and the sequence of x 's thus generated determines a polygonal line which approximates the continuous gradient ascent path joining the initial point in the sequence and the closest candidate local mode of the unknown probability density of the X 's. More precisely, the polygonal line generated by the repeated application of the update equation (1) from an initial point $x \in \mathbb{R}^q$ approximates the *integral curve* $\pi_x : \mathbb{R}_+ \rightarrow \mathbb{R}^q$ that solves the initial value problem

$$\begin{cases} \pi'_x(t) = \nabla p(\pi_x(t)) \\ \pi_x(0) = x, \end{cases} \quad (2)$$

where p is the unknown probability density. The curve π_x is also called a *gradient flow line*. A theoretical discussion about the connection between gradient ascent and the mean-shift algorithm, including general results on the rates of convergence of the polygonal line generated by gradient ascent to the corresponding gradient flow line, can be found in [Arias-Castro et al. \(2013\)](#).

[Chacón \(2012\)](#) provides a formal justification of clustering based on high density regions and local modes by means of Morse theory. The basic idea is that a *cluster* can be defined as the set of points x in the domain of a probability density (with non-degenerate critical points) whose gradient ascent paths π_x culminate at the same local mode, i.e. a cluster is the *basin of attraction* of a local mode (see also [Li et al., 2007a](#)). From the theory of differential equations, it is known that a curve π_x satisfying the initial value problem (2) exists and is unique under appropriate regularity assumptions on p . Moreover, because distinct integral curves can only intersect at the stationary points of p , the equivalence class of points whose gradient flow lines culminate at the same local mode naturally forms a cluster, and the collection of these equivalence classes is a partition of the support of the density function. The definition of cluster as the basin of attraction of a local mode is particularly appealing from an inferential viewpoint, because it endows cluster analysis with a sound theoretical foundation. By embracing this definition, we have a clear dichotomy between the notion of *population cluster*, which corresponds to the set of points associated to the high density region surrounding a local mode of the unknown probability density, and that of *empirical cluster*, corresponding to the set of points associated to the high density region surrounding a local mode of the estimated probability density.

[Hartigan \(1975\)](#) defines the population clusters of a probability distribution with density p as the connected components of the upper level sets of p . The connected components of the set $L(\lambda) = \{x \in \mathbb{R}^q : p(x) \geq \lambda\}$ for a given $\lambda \geq 0$, are considered the population clusters of p . These clusters can be estimated by $\hat{L}(\lambda) = \{x \in \mathbb{R}^q : \hat{p}(x) \geq \lambda\}$, as soon as an estimate \hat{p} of p is available, or otherwise by means of some other estimate of $L(\lambda)$ (see, among others, [Cuevas and Fraiman, 1997](#); [Stuetzle, 2003](#); [Azzalini and Torelli, 2007](#); [Stuetzle and Nugent, 2010](#); [Rinaldo and Wasserman, 2010](#)). However, the fact that this definition of population cluster depends on the resolution parameter λ is a severe limitation and may represent a drawback in certain applications (see, for instance, [Rinaldo et al., 2012](#)). Furthermore, the computation of plug-in estimates of $L(\lambda)$ and their connected components poses some difficulties and often requires intricate algorithms. Finally, a completely different approach to mode-based clustering based on persistent homology is considered by [Chazal et al. \(2013\)](#).

Although the problem of estimating the local modes of a probability density and using them for cluster analysis has been extensively studied in the Euclidean case, and the mean-

shift algorithm is widely used to perform this task, the same challenge has not received the same attention with *functional data*. Functional Data Analysis (Bosq, 2000; Ramsay and Silverman, 2005; Ferraty and Vieu, 2006; Horváth and Kokoszka, 2012) is a modern branch of statistics which deals with data that are intrinsically infinite dimensional such as functions, curves or surfaces. In the last two decades, several statistical methods (both parametric and nonparametric) that are fully developed for Euclidean data have been extended to the setting of random variables with infinite dimensional realizations. The theory of mode estimation for functional data is recent (see, for instance, Gasser et al., 1998; Daboniang et al., 2004 or Ferraty and Vieu, 2006 and references therein), but the estimation of *local* modes for functional data has received less attention, and a principled extension of the mean-shift algorithm that is adapted to this type of data has not been devised. Similarly, the literature on clustering methods for functional data is not as vast as that for Euclidean data (the reader can refer to Jacques and Preda, 2013 for a recent review). In particular, density-based clustering has received scant attention in the setting of functional data, mainly because of the difficulty related to defining a probability density in infinite dimensional spaces. In this paper, we address these challenges by providing a statistical framework that allows us to define and interpret a mean-shift algorithm for functional data, which we call the *functional mean-shift algorithm*.

3. The mean-shift algorithm for Euclidean data

The original mean-shift algorithm was proposed by Fukunaga and Hostetler (1975). The algorithm consists of repeatedly shifting a finite set of data points in the q dimensional Euclidean space towards their local sample means. More precisely, let $\mathcal{S} = \{X_1, \dots, X_n\} \subset \mathbb{R}^q$ be a sample of i.i.d. random vectors with distribution P that has density p with respect to the Lebesgue measure. Then, Fukunaga and Hostetler (1975) define the mean-shift operator $m(\cdot)$ as

$$m(x) = \frac{\sum_{X \in \mathcal{S}} K_h(X - x)X}{\sum_{X \in \mathcal{S}} K_h(X - x)} - x, \quad (3)$$

where K is a kernel function and $K_h(\cdot) = k(\|\cdot\|/h)$. For instance, K may correspond to the boxcar kernel

$$K(x) = k(\|x\|) = \begin{cases} 1 & \text{if } \|x\| \leq 1 \\ 0 & \text{if } \|x\| > 1 \end{cases},$$

or the Gaussian kernel

$$K(x) = k(\|x\|) \propto e^{-\frac{\|x\|^2}{2}}.$$

Here, $\|\cdot\|$ indicates the standard Euclidean norm, $h > 0$ is a bandwidth parameter and

$$x + m(x) = \frac{\sum_{X \in \mathcal{S}} K_h(X - x)X}{\sum_{X \in \mathcal{S}} K_h(X - x)}$$

is a local sample mean. For a general kernel function $K(x) = k(\|x\|)$, the associated function k is referred to as the *profile* of K .

The original mean-shift algorithm consists of repeatedly updating all the data points $X \in \mathcal{S}$ simultaneously according to the mean-shift update

$$X \leftarrow X + m(X) \quad (4)$$

until convergence. Because the data points $X \in \mathcal{S}$ tend to converge to a finite number of points $\tilde{\mathcal{M}} = \{\tilde{\mu}_1, \dots, \tilde{\mu}_r\} \subset \mathbb{R}^q$, the algorithm can be used to perform clustering. In particular, the j -th cluster of data points is defined as the subset of X 's in \mathcal{S} all converging to $\tilde{\mu}_j$ after the repeated iteration of the mean-shift update.

Cheng (1995) generalizes the original mean-shift algorithm proposed by Fukunaga and Hostetler (1975) and shows that the mean-shift algorithm includes k -means clustering as a limit case. He considers general isotropic kernels K and allows for the presence of positive weights ω in the local averaging of the sample points, so that the mean-shift operator of equation (3) now reads

$$m(x) = \frac{\sum_{X \in \mathcal{S}} \omega(X) K_h(X - x) X}{\sum_{X \in \mathcal{S}} \omega(X) K_h(X - x)} - x. \quad (5)$$

Most importantly, Cheng (1995) allows the mean-shift operator m to act on any arbitrary subset of points in \mathbb{R}^q while keeping the sample \mathcal{S} fixed, as opposed to the original Fukunaga and Hostetler (1975) version which applies equation (4) simultaneously to all the observed data points $X \in \mathcal{S}$, in fact updating the entire sample $\mathcal{S} \leftarrow \mathcal{S} + m(\mathcal{S})$ and thus generating a ‘blurring’ sequence of multi-sets. Today, Cheng’s version of the algorithm is called the mean-shift algorithm, and the original algorithm of Fukunaga and Hostetler is instead referred to as the blurring mean-shift algorithm.

When using Cheng’s mean-shift algorithm, one can imagine following the path generated by the update equation

$$x \leftarrow x + m(x) = \frac{\sum_{X \in \mathcal{S}} \omega(X) K_h(X - x) X}{\sum_{X \in \mathcal{S}} \omega(X) K_h(X - x)} \quad (6)$$

from any arbitrary initial $x \in \mathbb{R}^q$ while keeping the data points $X \in \mathcal{S}$ fixed. This operation corresponds to performing gradient ascent on the kernel density estimate

$$\tilde{p}(x) = \frac{1}{nh^q} \sum_{X \in \mathcal{S}} \omega(X) G_h(X - x)$$

of the true density p of the X ’s based on a second kernel G (often called the *shadow* of K). It is easy to verify that the mean-shift update of equation (6) corresponds to an update of the type

$$x \leftarrow x + s(x) \nabla \tilde{p}(x),$$

where $s(x)$ is a step size parameter depending on the current position x , and $\nabla \tilde{p}(x)$ corresponds to the gradient of \tilde{p} at x . Equivalently, one can rewrite the update above as

$$x \leftarrow x + s(x) \|\nabla \tilde{p}(x)\| a^*(x) = x + \bar{s}(x) a^*(x), \quad (7)$$

where the adaptive step size is

$$\bar{s}(x) \propto \|\nabla \tilde{p}(x)\| \frac{h^2}{\tilde{p}(x)}, \quad (8)$$

a^* is the unitary norm vector in the direction of the gradient $\nabla \tilde{p}(x)$ and

$$\hat{p}(x) = \frac{1}{nh^q} \sum_{X \in \mathcal{S}} \omega(X) K_h(X - x)$$

is a kernel density estimate of the unknown density p using the kernel K . The unit norm vector a^* gives the direction of steepest ascent at x . The important feature of equation (8) is the adaptive nature of the step size: if the current position x corresponds to a low density position (i.e. x is far from a local mode), then $\hat{p}(x)$ is small and the step in the direction of the gradient ascent is large; conversely, if the current position x corresponds to a high density position (i.e. x already is close to a local mode), then $\hat{p}(x)$ is large and the step in the direction of gradient ascent is small. Thus, thanks to the adaptivity of the step size, low density regions are only visited for a small number of iterations while, at the same time, the algorithm does not overshoot in high density regions.

The repeated application of the mean-shift update starting from an initial arbitrary point $x \in \mathbb{R}^q$ generates a trajectory that eventually converges to a (candidate) local mode of the estimated density \tilde{p} . This trajectory, which is a polygonal line, can be viewed as an estimate of the unknown gradient ascent path (or *gradient flow line*) $\pi_x : \mathbb{R}_+ \rightarrow \mathbb{R}^q$ that solves the initial value problem

$$\begin{cases} \pi'_x(t) = \nabla p(\pi_x(t)) \\ \pi_x(0) = x. \end{cases}$$

This feature of the algorithm allows us to determine a partition of \mathbb{R}^q in which each set of the partition corresponds to the basin of attraction of a local mode of the estimated density \tilde{p} , i.e. the equivalence class of all the points $x \in \mathbb{R}^q$ whose mean-shift trajectories culminate at the same local mode $\tilde{\mu}$ of \tilde{p} (which is an estimate of the equivalence class of points $x \in \mathbb{R}^q$ whose unknown gradient flow lines of p converge to the same unknown local mode μ of p). Similarly, a cluster of data points in the sample \mathcal{S} can be defined as a subset of X 's in \mathcal{S} whose mean-shift trajectories culminate at the same local mode of the estimated density \tilde{p} .

If the true density p was known, both the location and the number of the local modes $\mathcal{M} = \{\mu_1, \dots, \mu_R\}$ of p would be known. Therefore, if one accepts the above definition of cluster as basin of attraction of a local mode of the density, knowing p implies knowing the true *population* clustering of \mathbb{R}^q based on p . Because the true density p is generally unknown, the set \mathcal{M} (and therefore the true population clusters) are instead estimated by performing gradient ascent on an estimate \tilde{p} of the true density p . The estimated local modes $\tilde{\mathcal{M}} = \{\tilde{\mu}_1, \dots, \tilde{\mu}_r\}$ and their corresponding basins of attraction based on \tilde{p} thus provide a natural estimate of the unknown true population clustering. Recent results about the convergence of empirical clusters to their corresponding population clusters have been obtained by Chacón (2014).

4. A statistical framework for the functional mean-shift algorithm

In the remainder of the paper, we assume that the data $\mathcal{S} = \{X_1, \dots, X_n\}$ consist of n i.i.d. random variables defined on a suitable abstract probability space (Ω, \mathcal{A}, P) and we further assume that each random variable $X \in \mathcal{S}$ takes values in a (potentially) infinite dimensional measurable Hilbert space $(\mathcal{X}, \langle \cdot, \cdot \rangle)$. The inner product induces a norm $\|\cdot\|$ on \mathcal{X} , defined for any element $x \in \mathcal{X}$ as $\|x\| = \sqrt{\langle x, x \rangle}$. The norm $\|\cdot\|$, in turn, generates a distance d , defined between any pair x, y of elements of \mathcal{X} as $d(x, y) = \|x - y\|$.

Often it can be difficult to define a density for P if \mathcal{X} lacks a natural dominating measure. Whenever this is the case, we replace the notion of density with that of a *surrogate density*

as in [Gasser et al. \(1998\)](#) and [Ferraty et al. \(2012\)](#). Specifically, let $B(x, \epsilon) \subseteq \mathcal{X}$ indicate the open ball of radius ϵ centered at x in the topology induced by d and assume that there exist a function ϕ and a well-behaved functional p such that for any $x \in \mathcal{X}$

$$\varphi_x(\epsilon) = P(X \in B(x, \epsilon)) \sim p(x)\phi(\epsilon), \quad (9)$$

meaning that the so-called *small ball probability function* $\varphi_x(\epsilon)$ satisfies $\varphi_x(\epsilon) = p(x)\phi(\epsilon) + o(\phi(\epsilon))$ as $\epsilon \rightarrow 0$. Then, we consider the functional $p : \mathcal{X} \rightarrow \mathbb{R}_+$ satisfying the identifiability condition $E(p(X)) = 1$ a surrogate probability density of P on \mathcal{X} , and the function ϕ (often referred to as the *concentration function*) can be considered a volume parameter. [Ferraty et al. \(2006\)](#) and [Ferraty et al. \(2012\)](#) provide some examples of stochastic processes whose small ball probabilities fulfill the decomposition of equation (9). Notice that in the simpler setting in which X is a continuous real valued random variable with cumulative distribution function F , one has $\varphi_x(\epsilon) = P(X \in B(x, \epsilon)) = F(x + \epsilon) - F(x - \epsilon) = p(x)2\epsilon + o(\epsilon)$ where p here corresponds to the usual Radon-Nikodym derivative of P with respect to the Lebesgue measure and $\phi(\epsilon) = 2\epsilon$ is the Lebesgue measure of $B(x, \epsilon)$. However, for our purposes, we do not assume the existence of a dominating measure for P .

[Ferraty et al. \(2012\)](#) propose to estimate the population surrogate density p by means of

$$\hat{p}(x) = \frac{\frac{1}{n} \sum_{i=1}^n K_h(X_i, x)}{\frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i} K_h(X_i, X_j)} = w_K(\mathcal{S}) \sum_{X \in \mathcal{S}} K_h(X, x), \quad (10)$$

which is well-defined for $h \geq \min_{i,j} d(X_i, X_j)$. The kernel K is defined in terms of its profile k as

$$K(x, y) = k(d(x, y)), \quad (11)$$

and the profile k is a function of the distance d between two points x and y . The estimator is consistent under mild assumptions ([Ferraty et al., 2012](#), assumptions H1–H4). We henceforth assume that $k \geq 0$, k has compact support $[0, 1]$, $\int_{\mathbb{R}} k(t) dt = 1$ (although, as pointed out by [Cheng, 1995](#), the unit integral condition is superfluous for the mean-shift algorithm), there exist constants $-\infty < C_1 < C_2 \leq 0$ such that $C_2 \leq k' \leq C_1$ in $(0, 1)$ and k'' exists in $(0, 1)$.

It is worth mentioning that other surrogate densities have been proposed for random variables valued in infinite dimensional spaces. For instance, the reader may refer to [Delaigle and Hall \(2010\)](#) who define a surrogate probability density for function-valued random variables on the basis of the eigendecomposition obtained by means of principal component analysis.

Remark 1 *Assuming that the factorization of the small ball probability of equation (9) holds is not strictly necessary in practice, although it helps to establish a clearer connection between the standard mean-shift algorithm and the functional mean-shift algorithm in the rest of our discussion. When the factorization does not hold, one can work with the numerator of the estimator of equation (10),*

$$\frac{1}{n} \sum_{X \in \mathcal{S}} K_h(X, x)$$

and view it as an estimator of its own expected value, $\mathbb{E}_P(K_h(X, x))$. In turn, $\mathbb{E}_P(K_h(X, x))$ can be interpreted as a smooth population functional which (for some $h > 0$) is informative about local features of P .

Remark 2 *The choice of the distance in equation (11) should not be perceived as a complication, but rather as a useful element of flexibility in applied work. While on the one hand more traditional methods implicitly impose to the user the choice of the space in which to embed the data (e.g. functional PCA forces the user towards L_2 spaces and L_2 distances), the functional mean-shift algorithm allows the user to incorporate prior knowledge of the data generating process by choosing in which space to embed the data. This is implicitly done by tuning the distance function. In many practical problems, this flexibility can save a lot of work in pre-processing the data. For example, suppose that we observe a sample of curves that exhibit the same basic pattern except for a completely random and uninformative vertical shift, i.e. $X_i(t) = X_j(t) + \delta_{ij}$ where $\delta_{ij} \in \mathbb{R}$. In this case, a (semi-)distance based on the first derivatives such as*

$$d(x, y) = \|x' - y'\|_{L_2}$$

avoids the need to perform a vertical alignment of the curves before further analysis. In fact, in this case, $d(X_i(t), X_j(t)) = 0$. Similar considerations apply for systematic differences in the higher order derivatives among the curves (see [Ferraty and Vieu, 2006](#)).

5. The functional mean-shift algorithm

In this section, we show that applying the mean-shift algorithm to the elements of an Hilbert space using a kernel K is equivalent to applying an adaptive gradient ascent algorithm based on the Gateaux derivative of the estimated surrogate density

$$\tilde{p}(x) = w_G(\mathcal{S}) \sum_{X \in \mathcal{S}} G_h(X, x)$$

based on a second kernel, G . The kernel G , which is often referred to as the *shadow* of K , has a profile g which is related to the profile k of K in a particular way. We further discuss the notion of shadow of a kernel and the link between the profile k and g later in this section. We show that the gradient ascent direction can be characterized as the element of \mathcal{X} that has unit norm and maximizes the Gateaux derivative of \tilde{p} .

Before proceeding, let us briefly recall the definition of Gateaux differential and derivative.

Definition 3 (Gateaux differentiability, [Ambrosetti and Prodi, 1995](#)) *Let \mathcal{B} be a Banach space and let U be an open subset of \mathcal{B} . The map $F : U \mapsto V$ is Gateaux differentiable at $u \in U$ if there exists a continuous linear map F_u from \mathcal{B} to V , called the Gateaux differential of F at u , such that for all $w \in U$*

$$\lim_{\alpha \rightarrow 0} \frac{F(u + \alpha w) - F(u)}{\alpha} = F_u(w),$$

where $\alpha \in \mathbb{R}$. If F is Gateaux differentiable at all $u \in U$, then F is said to be Gateaux differentiable in U . The map $F'_G : u \mapsto F_u$, which associates to each $u \in U$ the continuous linear operator F_u , is called the Gateaux derivative of F .

The Gateaux derivative of the estimated surrogate density \tilde{p} is obtained in the following Lemma.

Lemma 4 (Gateaux differential of the estimated surrogate density) *Let $\tilde{p}_x(y)$ denote the Gateaux differential of \tilde{p} at $x \in \mathcal{X}$ evaluated at $y \in \mathcal{X}$. Let $h(X)$ be an adaptive bandwidth possibly varying across the data points $X \in \mathcal{S}$ which possibly depends on the entire sample \mathcal{S} , but not on x . Then,*

$$\tilde{p}_x(y) = \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x) (X - x) \right\rangle, \quad (12)$$

where the profiles g and k of the kernel G used in \tilde{p} and the kernel K satisfy $k(t) = -g'(t)/(Ct)$ for a constant $C > 0$.

Remark 5 *If the profile k of K is not continuous at 1 (as it is the case for the truncated Gaussian profile for instance), the Gateaux derivative of \tilde{p} is still almost surely continuous at any fixed $x \in \mathcal{X}$ as long as $P(\bigcup_{X \in \mathcal{S}} \{d(X, x) = h(X)\}) = 0$.*

In the second last line of the proof of the above Lemma (see Appendix), we set

$$k(t) = -\frac{g'(t)}{Ct} \quad (13)$$

with

$$C = \int_{\mathbb{R}} -\frac{g'(t)}{t} dt > 0.$$

One can verify that equation (13) (which is a differential version of equation (15) of [Cheng, 1995](#)) implies, among others, the correspondences of Table 1. For our purposes, we give the following definition of the shadow of a kernel.

Definition 6 (Shadow of a kernel) *If the profile k of the kernel K used in the mean-shift algorithm and the profile g of the kernel G used in the corresponding gradient ascent algorithm satisfy equation (13), then G is said to be the shadow of K .*

As a matter of fact, Gaussian kernels are the only kernels that coincide with their own shadows; see Table 1 and ([Cheng, 1995](#), Theorem 2).

When the Gateaux derivative of a functional $F : \mathcal{X} \rightarrow \mathbb{R}$ at a point $x \in \mathcal{X}$ can be expressed as an inner product of the form $\langle y, \nabla \tilde{p}_x \rangle$ for any $y \in \mathcal{X}$, then the element $\nabla \tilde{p}_x \in \mathcal{X}$ is the *functional gradient of \tilde{p} at x* . Thus, the function

$$\nabla \tilde{p}_x = Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x) (X - x) \quad (14)$$

of equation (12) is the functional gradient of \tilde{p} at x .

We can rewrite equation (14) as

$$\nabla \tilde{p}_x = Cw_G(\mathcal{S}) m(x) \bar{p}(x), \quad (15)$$

where \bar{p} is a weighted and unnormalized estimate of the surrogate density p based on K ,

$$\bar{p}(x) = \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x),$$

Table 1: Some correspondences between the profile k of the kernel K and the profile g of its shadow G implied by equation (13).

G	K
$g(t) \propto (1 - t^2) \mathbb{1}_{[0,1]}(t)$ (Epanechnikov)	$k(t) \propto \mathbb{1}_{[0,1]}(t)$ (uniform)
$g(t) \propto (1 - t^2)^2 \mathbb{1}_{[0,1]}(t)$ (biweight)	$k(t) \propto (1 - t^2) \mathbb{1}_{[0,1]}(t)$ (Epanechnikov)
$g(t) \propto (1 - t^2)^3 \mathbb{1}_{[0,1]}(t)$ (triweight)	$k(t) \propto (1 - t^2)^2 \mathbb{1}_{[0,1]}(t)$ (biweight)
$g(t) \propto \cos\left(\frac{\pi}{2}t\right) \mathbb{1}_{[0,1]}(t)$ (cosine)	$k(t) \propto \frac{\sin(\frac{\pi}{2}t)}{t} \mathbb{1}_{[0,1]}(t)$ (sinc)
$g(t) \propto e^{-\frac{t^2}{2}} \mathbb{1}_{[0,1]}(t)$ (Gaussian)	$k(t) \propto e^{-\frac{t^2}{2}} \mathbb{1}_{[0,1]}(t)$ (Gaussian)

and m is the functional mean-shift operator

$$m(x) = \frac{\sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x) X}{\sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x)} - x. \quad (16)$$

Notice that equation (16) is a functional analog of equation (5) with $\omega(X) = 1/h^2(X)$. Let $a^*(x)$ denote the element with unit norm in the direction of $\nabla \tilde{p}_x$ from x , i.e.

$$a^*(x) = \frac{\nabla \tilde{p}_x}{\|\nabla \tilde{p}_x\|}.$$

Then, from (15) it follows that

$$m(x) = \frac{\sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x) X}{\sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x)} - x = \frac{\|\nabla \tilde{p}_x\|}{C w_G(\mathcal{S}) \bar{p}(x)} a^*(x), \quad (17)$$

hence the functional mean-shift update

$$x \leftarrow x + m(x) = \frac{\sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x) X}{\sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x)} \quad (18)$$

corresponds to a gradient ascent update $x \leftarrow x + s(x) a^*(x)$ that is conceptually identical to equation (7), with a step size

$$s(x) = \frac{\|\nabla \tilde{p}_x\|}{C w_G(\mathcal{S}) \bar{p}(x)} \quad (19)$$

that is conceptually identical to (8). Notice further that if $x^* \in \mathcal{X}$ is a fixed point of the functional mean-shift update of equation (18) (i.e. $m(x^*) = 0$ in equation (17)), then it is easily seen that $\nabla \tilde{p}_{x^*} = 0$.

It is now clear that, because the functional mean-shift algorithm is a gradient ascent algorithm in an infinite-dimensional space, the trajectory generated by the repeated application of the functional mean-shift operator of equation (18) can be thought of as an

estimate of the gradient flow line associated to the initial value problem

$$\begin{cases} \pi'_x(t) = \nabla p_{\pi_x(t)} \\ \pi_x(0) = x, \end{cases}$$

whenever such flow exists. It is known from the theory of differential equations with Lipschitz coefficients that, in a generic Hilbert space \mathcal{X} , the initial value problem above has exactly one solution for any starting point $x \in \mathcal{X}$ if functional gradient of p is a Lipschitz map, i.e. if there exists $L > 0$ such that for any $x, y \in \mathcal{X}$ we have $\|\nabla p_x - \nabla p_y\| \leq L\|x - y\|$. It is easy to check that the functional gradient of the estimated surrogate density is a Lipschitz map, therefore the functional gradient flow starting from an arbitrary point $x \in \mathcal{X}$ exists and it is unique on the estimated surrogate density. Furthermore, because any finite sample only spans a finite dimensional subspace of \mathcal{X} , the gradient flows on the estimated surrogate density also converge. The Lipschitz condition, however, is not enough to guarantee the convergence of the gradient flows as $t \rightarrow \infty$ in general, and in particular it does not guarantee the convergence of the gradient flows of a general population surrogate density p . Convergence, indeed, is a delicate question (especially in infinite dimensional spaces) and usually requires strong regularity conditions on p . We do not address this question in this paper, nor do we address the question of the convergence of the sequences generated by the mean-shift algorithm (which is a topic of active research even in finite dimensions; see, among others, [Li et al., 2007b](#) and [Aliyari Ghassabeh, 2013](#)). Intuitively, however, a sequence of elements of \mathcal{X} obtained by means of the functional mean-shift algorithm approximates the corresponding gradient flow on the estimated surrogate density \tilde{p} , and can be thought of as an estimate of the associated gradient flow on the unknown population surrogate density p (whenever such flow exists).

Remark 7 *If the bandwidth is fixed rather than adaptive, i.e. $h(X) = h > 0$ for all $X \in \mathcal{S}$, then the step size of equation (19) is*

$$s(x) = \frac{\|\nabla \tilde{p}_x\|}{Cw_G(\mathcal{S}) \bar{\bar{p}}(x)} \frac{h^2}{\bar{\bar{p}}(x)} \propto \|\nabla \tilde{p}_x\| \frac{h^2}{\bar{\bar{p}}(x)},$$

where now the fixed bandwidth is not incorporated in the unnormalized surrogate density estimate

$$\bar{\bar{p}}(x) = \sum_{X \in \mathcal{S}} K_h(X, x),$$

and the analogy with equation (8) is even more evident.

Remark 8 *The first order condition*

$$Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x) (X - x) = 0$$

is a necessary optimality condition that is satisfied by all critical points of the surrogate density estimate \tilde{p} , and by the uninteresting trivial roots in the set $\mathcal{X} \setminus \bigcup_{X \in \mathcal{S}} \bar{B}(X, h(X))$, where $\bar{B}(X, h(X))$ denotes the closure of the open ball. This condition is not sufficient for local maxima and thus every $x \in \bigcup_{X \in \mathcal{S}} \bar{B}(X, h(X))$ satisfying $\nabla \tilde{p}_x = 0$ should be considered

a candidate local mode of the estimated surrogate density \tilde{p} . To verify that a solution $x^* \in \bigcup_{X \in \mathcal{S}} \bar{B}(X, h(X))$ of the functional equation $\nabla \tilde{p}_x = 0$ is in fact a local mode and rule out unstable maxima (such as saddle points or plateaus), one can apply a suitable perturbation δ to x^* and then reapply the functional mean-shift algorithm to $x^* + \delta$. If the sequence of functions generated by the functional mean-shift algorithm starting from $x^* + \delta$ still converges to x^* , then x^* can be regarded as a local mode of the estimated surrogate density.

In the Euclidean setting, [Genovese et al. \(2013\)](#) develop a procedure to test the hypothesis $H_0 : x^*$ is not a mode of the population density p versus $H_1 : x^*$ is a mode of the population density p . We provide a similar test for the functional case in the next section.

Remark 9 Equation (13) imposes some restrictions on the choice of the profiles k and g . In fact, once we fix the profile g in the gradient ascent scheme, the profile k of the corresponding mean-shift algorithm must satisfy the differential inequality

$$k'(t) = \frac{tg''(t) - g'(t)}{t^2} \leq 0 \quad (20)$$

for $t \in (0, 1)$ and the condition

$$\lim_{t \rightarrow 0^+} k(t) = \lim_{t \rightarrow 0^+} -\frac{g'(t)}{t} = c \quad (21)$$

with $0 < c < \infty$. In particular, (20) implies that g must satisfy

$$g(t) \leq c_1 + c_2 t^2$$

for some $c_1 > 0$ (to ensure that g is positive at $t = 0$ and $c_2 < 0$ (to ensure that g is decreasing for $t \geq 0$), while (21) requires that g is locally quadratic in a (right) neighborhood of $t = 0$. All the profiles presented in Table 1 satisfy these two conditions.

Remark 10 Another useful property of the mean-shift algorithm is that it automatically detects outliers in the sample \mathcal{S} . If an element $X^* \in \mathcal{S}$ is separated from the other elements $X \in \mathcal{S}$ with respect to the norm $\|\cdot\|$ of the Hilbert space in which the data are embedded, then the algorithm generates the atomic cluster $\{X^*\}$. To illustrate, suppose that we are able to choose the bandwidths $h(X)$ according to a ‘good’ procedure. Then, if

$$X^* \notin \bigcup_{X \in \mathcal{S}} \bar{B}(X, h(X)),$$

we have $\nabla \tilde{p}_{X^*} = C w_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, X^*) (X - X^*) = 0$ as long as K has a profile that is supported on $[0, 1]$. Atomic clusters produced by the functional mean-shift algorithm under a ‘good’ choice of the bandwidth should therefore be regarded as potential outliers.

We conclude this section by noting that, in analogy with the Euclidean case, the blurring version of the functional mean-shift algorithm is easily obtained by iterating

$$X \leftarrow X + m(X)$$

on all $X \in \mathcal{S}$ simultaneously, with m as in equation (16).

6. Connection with clustering using fPCA and k -means

For the clustering of a sample of functional data that have a small intrinsic dimensionality (say, D), the functional mean-shift algorithm under the L_2 distance is expected to perform at least as good as applying k -means clustering to the projection coefficients on the first D principal components. If the intrinsic dimensionality is in fact D , then the squared distance between two functional data points is

$$d^2(X_i, X_j) = \|X_i - X_j\|^2 \approx \sum_{k=1}^D (\theta_{i,k} - \theta_{j,k})^2,$$

where $\theta_{i,k}$ is the projection coefficient of X_i on the k -th principal component. In this case, clustering using the functional mean-shift algorithm is essentially equivalent to density-based clustering using a kernel density estimate based on the D -dimensional projection coefficients. k -means tends to produce accurate clusterings when the number of clusters is known and the empirical clusters have elliptical shapes. Hence, roughly speaking, if D is small, the number of clusters is known, and the clusters of principal component scores have elliptical shapes, then the functional mean-shift algorithm with the L_2 distance is expected to yield very similar results when compared to k -means clustering on the D -dimensional projection coefficients of fPCA (Figure 3). However, when the intrinsic dimensionality of the functional data is large, the functional mean-shift algorithm has greater flexibility compared to the combined fPCA/ k -means clustering, in that it can pick up differences between the clusters that may not be evident in the projections onto the lower dimensional space spanned by the first few principal components. Also, as opposed to the fPCA/ k -means approach, the functional mean-shift algorithm does not require the user to select the number of clusters a priori.

Interestingly, even in finite dimensional settings there exist situations in which the functional mean-shift algorithm outperforms the combined fPCA/ k -means analysis. For example, Figure 4 depicts the projections of a set of 200 curves that lie in the span of the functions $\sin(2\pi t)$ and $\cos(2\pi t)$ onto their first two principal components (which explain 100% of the sample variance). Here, k -means is doomed to fail with the circular pattern of Figure 4 even when the number of clusters is known and the seeds are placed suitably. Figure 5 depicts the clusters obtained by means of the combined fPCA/ k -means approach. It is clear that these clusters do not represent the underlying structure of the data. On the other hand, the functional mean-shift algorithm generates meaningful clusters (Figure 6) and effectively summarizes each of them by its modal curve.

7. Inference on the functional modes

In the previous section, we showed that the functional mean-shift algorithm locates the candidate local modes of the estimated surrogate density \tilde{p} , namely a set of points $\tilde{\mathcal{M}} = \{\tilde{\mu}_1, \dots, \tilde{\mu}_r\}$ such that the first order condition

$$\nabla \tilde{p}_{\mu_j} = 0 \iff \tilde{p}_{\mu_j}(y) = 0 \quad \forall y \in \mathcal{X}$$

is satisfied for $j = 1, \dots, r$. We may want to test whether $\tilde{\mu}_j \in \tilde{\mathcal{M}}$ is a statistically *significant* functional mode of the unknown surrogate density p , i.e. we may want to test

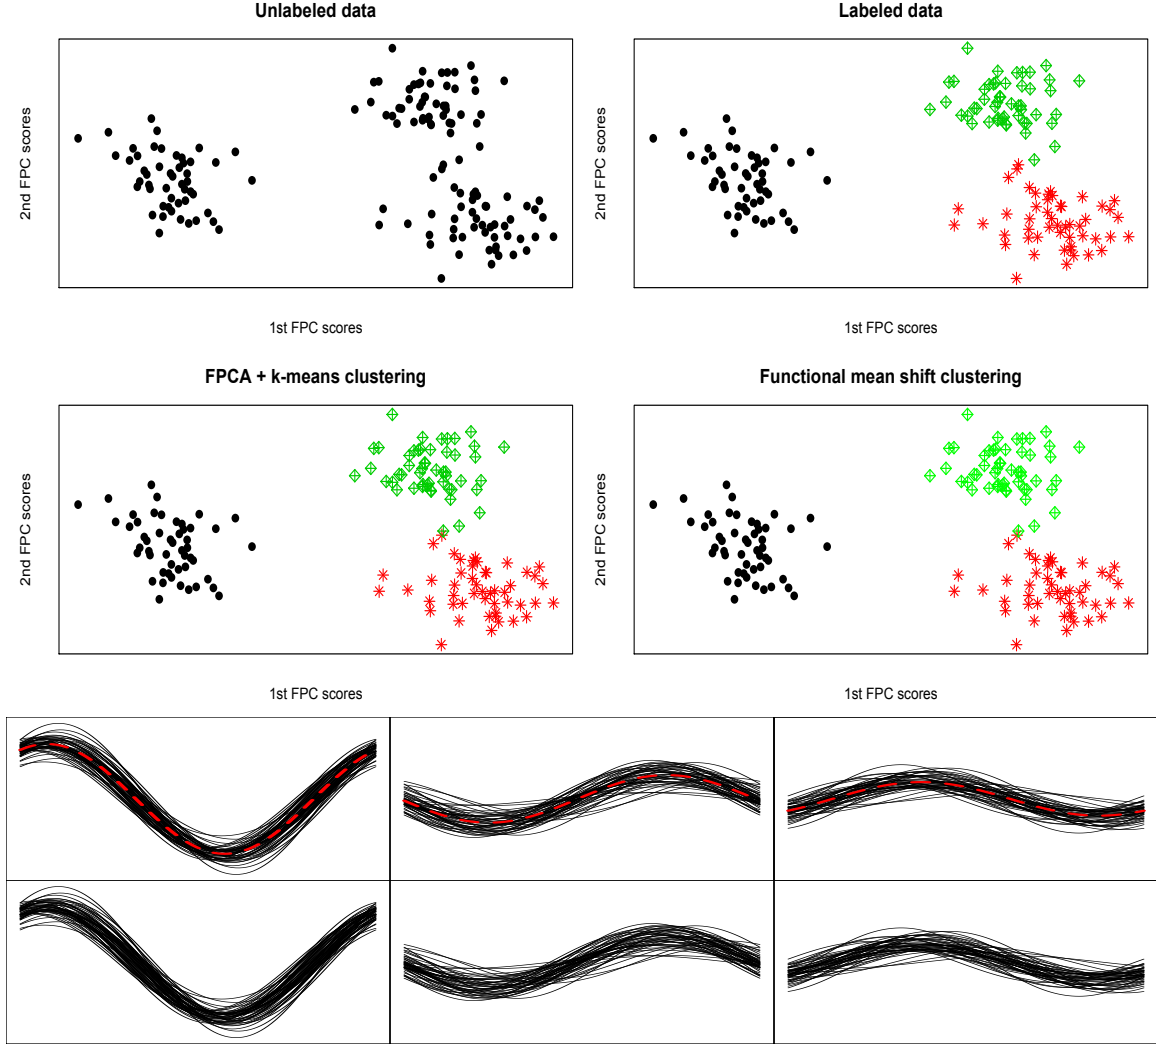


Figure 3: Top 4 panels: Projection coefficients on the first two principal components of a set of 150 curves in the span of $\sin(2\pi t)$ and $\cos(2\pi t)$. Bottom 4 panels: Clusters recovered by the functional mean-shift algorithm (top) with modal curves (red dashed lines) and clusters recovered by fPCA/ k -means (bottom). When the clusters have elliptical shapes, the combined fPCA/ k -means analysis and the functional mean-shift algorithm both recover the underlying clustering structure of the data.

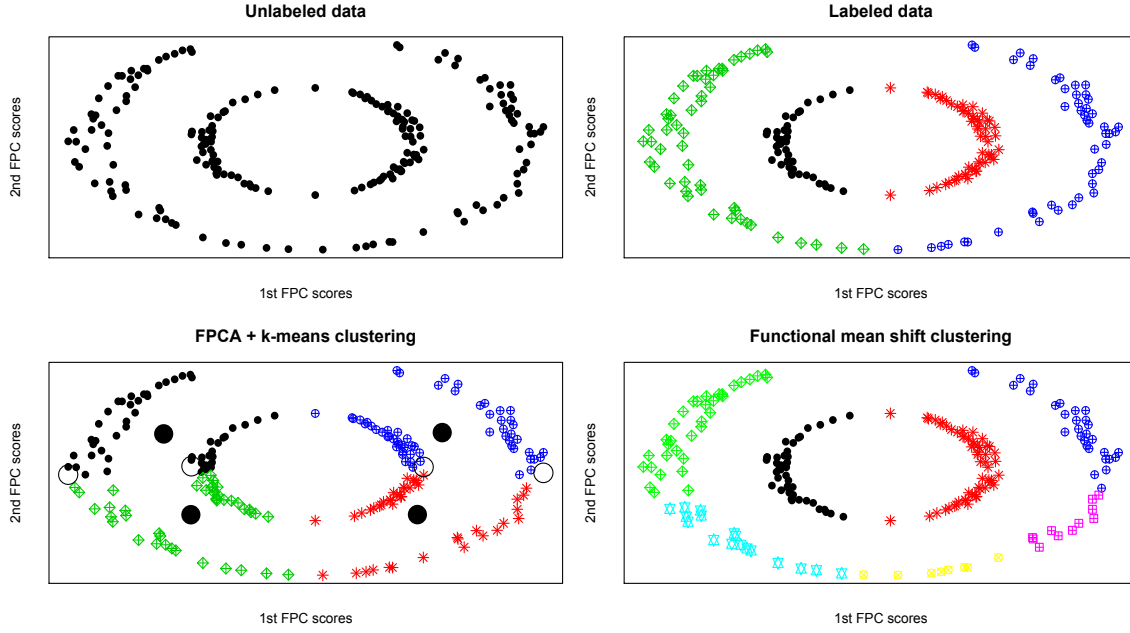


Figure 4: Projection coefficients on the first two principal components of a set of 200 curves in the span of $\sin(2\pi t)$ and $\cos(2\pi t)$. The circular pattern prevents k -means clustering from achieving a satisfactory result even when the number of clusters is known a priori and the initial seeds (hollow circles in the bottom left panel) are suitably placed. The large full circles in the bottom left panel represent the final cluster centers obtained with k -means.




Figure 5: Clusters obtained with the combined PCA/ k -means analysis.




Figure 6: Clusters obtained with the functional mean-shift algorithm. The modal curves are represented by the red dashed lines.

whether $\tilde{\mu}_j \in \tilde{\mathcal{M}}$ is a critical point of p and the second Gateaux derivative of p is a negative definite operator at $\tilde{\mu}_j$. Let us assume that the unknown surrogate density is twice Gateaux differentiable with second Gateaux differential at $x \in \mathcal{X}$ denoted by

$$p_x^{(2)} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R} \\ (y, z) \mapsto p_x^{(2)}(y, z),$$

which is a symmetric continuous bilinear operator (the analog of the Hessian matrix in infinite dimensions). A critical point $\mu_j \in \mathcal{M}$ of the unknown surrogate density p is a local maximum if and only if $p_{\mu_j}^{(2)}$ is negative definite, i.e.

$$\sup_{\|y\|=1} p_{\mu_j}^{(2)}(y, y) = \lambda_{\mu_j} < 0.$$

Thus, for $\tilde{\mu}_j \in \tilde{\mathcal{M}}$, we want to test

$$H_0 : \lambda_{\tilde{\mu}_j} \geq 0 \quad \text{or} \quad \nabla p_{\tilde{\mu}_j} \neq 0 \\ \text{vs.} \\ H_1 : \lambda_{\tilde{\mu}_j} < 0 \quad \text{and} \quad \nabla p_{\tilde{\mu}_j} = 0,$$

that is

$$H_0 : \tilde{\mu}_j \text{ is not a mode of } p \\ \text{vs.} \\ H_1 : \tilde{\mu}_j \text{ is a mode of } p.$$

The pointwise nature of the alternative hypothesis on the functional gradient may constitute a problem in the specification of the above hypothesis test. For instance, in the Euclidean

setting, the condition $\nabla p_{\tilde{\mu}_j} = 0$ would correspond to a zero (Lebesgue) measure set. At least in the Euclidean setting, one possibility is to recast the hypothesis test as follows (Genovese et al., 2013):

$$\begin{aligned} H_0 : \neg H_1 \\ \text{vs.} \\ H_1 : \lambda_{\tilde{\mu}_j} < 0 \quad \text{and} \quad \|\nabla p_{\tilde{\mu}_j}\| < \delta \end{aligned} \tag{22}$$

for some $\delta > 0$, so that we are instead testing the hypothesis that $\tilde{\mu}_j$ is an *approximate* mode of p . The sample splitting procedure introduced in the test of Genovese et al. (2013) is such that constraint on the gradient has essentially no effect in (22) (we will later return to this point in Remark 14). In the following, we consider a similar test that is adapted to the functional case. We focus on

$$\begin{aligned} H_0 : \lambda_{\tilde{\mu}_j} \geq 0 \\ \text{vs.} \\ H_1 : \lambda_{\tilde{\mu}_j} < 0. \end{aligned} \tag{23}$$

A natural test statistic for (23) is $\tilde{\lambda}_{\tilde{\mu}_j} = \sup_{\|y\|=1} \tilde{p}_{\tilde{\mu}_j}^{(2)}(y, y)$. Lemma 11 gives the explicit form of the second Gateaux differential of the estimated surrogate density at an arbitrary point $x \in \mathcal{X}$. The explicit form of the test statistic $\tilde{\lambda}_x$ is then derived in Lemma 12.

Lemma 11 (Second Gateaux differential of the estimated surrogate density) *The second Gateaux differential of the estimated surrogate density \tilde{p} at $x \in \mathcal{X}$ evaluated at $(y, z) \in \mathcal{X} \times \mathcal{X}$ is*

$$\begin{aligned} \tilde{p}_x^{(2)}(y, z) = -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) \frac{\langle X - x, y \rangle \langle X - x, z \rangle}{\|X - x\|} \right. \\ \left. + K_h(X, x) \langle y, z \rangle \right]. \end{aligned}$$

Lemma 12 (Test statistic for the second order optimality condition) *We have*

$$\begin{aligned} \tilde{\lambda}_x = \sup_{\|y\|=1} \tilde{p}^{(2)}(y, y) = Cw_G(\mathcal{S}) \left[2 \left\| \sum_{X \in \mathcal{S}} \frac{1}{h^3(X)} K'_h(X, x) \frac{X - x}{\|X - x\|} \right\| - \right. \\ \left. \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) (\|X - x\| + \|X - x\|^{-1}) + K_h(X, x) \right] \right]. \end{aligned}$$

Remark 13 *Notice that when the profiles g and k satisfy the conditions of Remark 9, $\lim_{t \rightarrow 0^+} \frac{k'(t)}{t}$ exists and it is finite. Thus, we can safely set $\frac{k'(0)}{0} = \lim_{t \rightarrow 0^+} \frac{k'(t)}{t}$ in the expressions involved in the two previous Lemmata whenever needed.*

We are now in the position to describe our testing procedure. Suppose for simplicity that the sample \mathcal{S} contains an even number of elements, say $2n$, and that we use a non-adaptive

bandwidth parameter $h(X) = h > 0$ for all $X \in \mathcal{S}$.

Stage 1: First, we divide the sample in two subsamples of size n . We apply the functional mean-shift algorithm on the first subsample, denoted X_1, \dots, X_n , in order to determine the set $\tilde{\mathcal{M}} = \{\tilde{\mu}_1, \dots, \tilde{\mu}_r\}$ of candidate local modes of \tilde{p} , which is an estimate of the set $\mathcal{M} = \{\mu_1, \dots, \mu_R\}$ of local modes of the unknown surrogate density p . This step usually requires using the first subsample X_1, \dots, X_n to determine the value of the bandwidth h for the functional mean-shift algorithm.

Stage 2: Next, we use the second subsample, denoted Y_1, \dots, Y_n , to compute the test statistic $\tilde{\lambda}_{\tilde{\mu}_j}$ for $j = 1, \dots, r$. For $b = 1, \dots, B$ we sample with replacement n elements of the second subsample, $Y_{1,b}^*, \dots, Y_{n,b}^*$ from the empirical distribution of Y_1, \dots, Y_n , $P_n(y) = \frac{1}{n} \sum_{i=1}^n \delta_{Y_i}(y)$. For $b = 1, \dots, B$ we compute the test statistic $\tilde{\lambda}_{\tilde{\mu}_j,b}^*$ using the bootstrap sample $Y_{1,b}^*, \dots, Y_{n,b}^*$ and we thus construct a $1 - \alpha$ bootstrap confidence interval for $\lambda_{\tilde{\mu}_j}$, $\mathcal{C}_{n,1-\alpha}(\lambda_{\tilde{\mu}_j})$ where $\alpha \in (0, 1)$. Finally, if $\sup \{\lambda \in \mathcal{C}_{n,1-\alpha}(\lambda_{\tilde{\mu}_j})\} < 0$, we reject the null hypothesis that $\tilde{\mu}_j$ is not a local mode of the unknown surrogate density p .

By setting $\alpha' = \alpha/\ell$, we can use the same procedure to construct a confidence rectangle for $\ell \in \{1, \dots, r\}$ of the $\lambda_{\tilde{\mu}_j}$'s simultaneously. This rectangle takes the form $\mathcal{C}_{n,1-\alpha}(\lambda_{\tilde{\mu}_{j_1}}, \dots, \lambda_{\tilde{\mu}_{j_\ell}}) = \mathcal{C}_{n,1-\alpha'}(\lambda_{\tilde{\mu}_{j_1}}) \times \dots \times \mathcal{C}_{n,1-\alpha'}(\lambda_{\tilde{\mu}_{j_\ell}})$. Note that the purpose of sample splitting in Stage 1 is to assure the validity of the confidence intervals. The simultaneous test algorithm is summarized in Figure 7.

Remark 14 *Because each candidate local mode $\tilde{\mu}_j \in \tilde{\mathcal{M}}$ is such that $\nabla \tilde{p}_{\tilde{\mu}_j} = 0$ in the first stage, the estimated gradient is likely to be null also in the second stage. Thus, the constraint $\|\nabla p_{\tilde{\mu}_j}\| < \delta$ has a negligible effect in (22).*

Let us now illustrate the above testing procedure by means of a simple simulation. We randomly draw 150 curves from these three groups:

- $X(t) = \eta_X \cos(5\pi/2 \cdot t)$
- $Y(t) = 3 + \eta_Y \cos(5\pi/2 \cdot t)$
- $C(t) = \gamma + 3 \cdot \mathbb{1}_{\{U > 0.5\}} + \cos(5\pi/2 \cdot t)$

where $\eta_X \stackrel{d}{=} \eta_Y \sim \mathcal{N}(\mu = 1, \sigma = 0.1)$, $\gamma \sim \mathcal{N}(\mu = 0, \sigma = 0.8)$, $U \sim \text{Unif}(0, 1)$, $t \in [0, 1]$, and all the random variables involved are independent of each other. The probability of drawing a curve from a given group is the same for each group and equal to $1/3$. In this simulation, the X 's and the Y 's are thought of as realizations of two distinct *signals*, whereas the C 's are *clutter* curves that correspond to a version of these signals that is subject to a vertical perturbation. On the basis of the data generating process, it makes sense to expect the existence of a modal curve μ_X for the X 's and a distinct modal curve μ_Y for the Y 's. However, the clutter curves are likely to generate some spurious local mode (at least for small enough bandwidths) when one estimates μ_X and μ_Y using the functional mean-shift algorithm. The simulated curves are presented in the top panel of Figure 8 (the X 's is grey, the Y 's in red and the C 's in light blue). The functional mean-shift identifies 3 candidate

Local functional mode set testing algorithm

Input: sample \mathcal{S} with $\#(\mathcal{S}) = 2n$, coverage level $1 - \alpha$

Output: a set $\tilde{\mathcal{M}}^S \subset \mathcal{X}$ of significant local modes at the approximate level $1 - \alpha$

1. split the sample in two halves, X_1, \dots, X_n and Y_1, \dots, Y_n
2. apply the mean-shift algorithm on X_1, \dots, X_n to find the set of candidate local modes $\tilde{\mathcal{M}} = \{\tilde{\mu}_1, \dots, \tilde{\mu}_r\}$
3. for each $j = 1, \dots, r$ and for $b = 1, \dots, B$ do
 - resample with replacement n elements $Y_{1,b}^*, \dots, Y_{n,b}^*$
 - compute $\tilde{\lambda}_{\tilde{\mu}_j, b}^*$
4. for each $j = 1, \dots, r$ construct the $1 - \alpha/r$ level confidence interval $\mathcal{C}_{n, 1 - \alpha/r}(\lambda_{\tilde{\mu}_j})$ using the empirical quantiles of $\tilde{\lambda}_{\tilde{\mu}_j, 1}^*, \dots, \tilde{\lambda}_{\tilde{\mu}_j, B}^*$.
5. for each $j = 1, \dots, r$ check the condition $\sup \{\lambda \in \mathcal{C}_{n, 1 - \alpha/r}(\lambda_{\tilde{\mu}_j})\} < 0$
6. set $\tilde{\mathcal{M}}^S = \left\{ \tilde{\mu}_j \in \tilde{\mathcal{M}} : \sup \{\lambda \in \mathcal{C}_{n, 1 - \alpha/r}(\lambda_{\tilde{\mu}_j})\} < 0 \right\}$
7. output $\tilde{\mathcal{M}}^S$.

Figure 7: The local functional mode set testing algorithm. Sample spitting assures the validity of the test.

functional modal curves $\tilde{\mu}_X$, $\tilde{\mu}_C$, and $\tilde{\mu}_Y$, with the asymmetric truncated Gaussian kernel (top panel of Figure 8). Inside the kernel, we set d to be the L_2 distance and the bandwidth is set to $h = 7.5$ (approximately the 41th percentile of the L_2 distances among the curves in the first subsample). As expected, only the local modes corresponding to the X 's and the Y 's are significant on the basis of the local mode testing algorithm of Figure 7 (in this simulation, we set the number of bootstrap replications B equal to 1000). The output of the testing procedure and the 95% simultaneous approximate confidence intervals for $\lambda_{\tilde{\mu}_X}$, $\lambda_{\tilde{\mu}_C}$, and $\lambda_{\tilde{\mu}_Y}$ are summarized in the bottom panel of Figure 8.

8. Application: spike sorting

In this section, we describe how the functional mean-shift algorithm can be used to cluster a set of curves that correspond to neural activity.

The available data, displayed in Figure 9, represent a subset of 40 recordings of neurons over time which come from a behavioral experiment performed at the Andrew Schwartz motorlab (<http://motorlab.neurobio.pitt.edu/index.php>) on a macaque monkey (the authors thank Andrew Schwartz, Valérie Ventura and Sonia Todorova for providing the data). The monkey performs a center-out and out-center target reaching task with 26 targets in a virtual 3D environment. The curves of Figure 9 show the voltage of neurons versus the times recorded at electrodes (32 equidistant time points per curve; time is standardized between 0 and 1). The recorded neural activity consists of all the action potentials detected above a channel-specific threshold on a 96-channel Utah array implanted in the primary motor cortex.

The goal is to perform *spike sorting*, i.e. clustering the curves in a set of distinct homogeneous groups. Each cluster of curves is then thought to correspond to the activity of a single neuron since each neuron tends to have its own characteristic curve (or *spike*). An analysis of (a larger set of) these curves can also be found in Lei et al. (2013), who cluster the curves following the *conformal prediction* approach (Vovk et al., 2009).

On the basis of Figure 9, one may expect to find 3 clusters, each summarized by a corresponding modal curve. We set d to be the distance associated with the Sobolev space H_1 of square integrable functions on the standard unit interval with square integrable first weak derivative, i.e.

$$d(x, y) = \|x - y\|_{L_2} + \|x' - y'\|_{L_2}.$$

This distance is associated to the inner product

$$\langle x, y \rangle_{H_1} = \langle x, y \rangle_{L_2} + \langle x', y' \rangle_{L_2}.$$

The curves are smoothed using local quadratic polynomials and their first derivatives are obtained directly from the local polynomial fit. We apply the functional mean-shift on this dataset of curves using the asymmetric truncated Gaussian kernel. A fixed bandwidth parameter is chosen using a heuristic based on the plot of the number of non-atomic clusters (that is, clusters of sample curves containing more than a single curve) as a function of bandwidth (Figure 10). In particular, we repeatedly run the functional mean-shift using an increasing sequence of bandwidths (100 equally separated values between 5% and 50% of the largest observed distance between the sample curves). The candidate bandwidths are chosen




Figure 8: Top panel: Simulated data. The red and the black curves (top and bottom bundles of continuous lines) are representative of two similar, but distinct, signals. The blue curves (dashed lines) are clutter curves. The highlighted curves are the 3 candidate functional local modes identified by the functional mean-shift algorithm, $\tilde{\mu}_X$, $\tilde{\mu}_C$, and $\tilde{\mu}_Y$. Bottom panel: Simultaneous approximate 95% confidence intervals for $\lambda_{\tilde{\mu}_X}$, $\lambda_{\tilde{\mu}_C}$, and $\lambda_{\tilde{\mu}_Y}$ using $B = 1000$ bootstrap replications. Only the confidence intervals associated to the candidate modal curves $\tilde{\mu}_X$ and $\tilde{\mu}_Y$ entirely lie to the left of 0. Therefore, as expected, only $\tilde{\mu}_X$ and $\tilde{\mu}_Y$ appear to correspond to significant functional local modes.




Figure 9: Plot of the 40 neural activity curves. Three curves are highlighted to show typical trajectories. On the basis of this figure, one expects that the curves can be clustered in 3 homogeneous groups, each characterized by its own modal trajectory.

as the midpoints of the ranges where the number of clusters stabilizes. For these data, this heuristic identifies $h_1 = 0.2150 \cdot \max_{X, X' \in \mathcal{S}} d(X, X')$ and $h_2 = 0.3525 \cdot \max_{X, X' \in \mathcal{S}} d(X, X')$ as the two candidate bandwidths. The output of the mean-shift algorithm is displayed in Figures 11 and 12. In this ‘multi-bandwidth’ analysis, the largest candidate bandwidth, h_2 , produces a coarse clustering of the sample curves: from Figure 11, we see that the two most similar bundles of curves are clustered together in a unique cluster. At this resolution, the algorithm recognizes that these two bundles have a more similar shape, which is different from the sigmoid shape of the third bundle. The smallest candidate bandwidth, h_1 , produces a clustering of the sample curves that is just right (Figure 12). At this resolution, the functional mean-shift algorithm also identifies two atomic clusters (dashed lines in the lower two panels of Figure 12) which, in light of Remark 10, could be considered potential outliers.

9. Application: signature forgery

Suppose that a dataset of n similar signatures, $\mathcal{S} = \{X_1, \dots, X_n\}$, is available to us. Each observed signature $X \in \mathcal{S}$ can be thought of as a (discretized version of a) curve in \mathbb{R}^2 of the type

$$t \mapsto X(t) = (x(t), y(t)),$$

where, for simplicity, we can take $t \in [0, 1]$. To simplify, we may identify each signature in the dataset with its *tangential acceleration*, i.e. the real valued function

$$t \mapsto S(t) = (x''(t), y''(t)) \cdot \frac{(x'(t), y'(t))}{\|(x'(t), y'(t))\|_2}, \quad (24)$$




Figure 10: Number of non-atomic clusters of neural activity curves as a function of the bandwidth parameter, h . The number of non-atomic clusters stabilizes in the 2 highlighted ranges. The candidate bandwidth parameters are then set to be the midpoint of the 2 relevant ranges. The numbers in parentheses above the candidate bandwidths denote the number of sample curves that are clustered in non-atomic clusters using the corresponding candidate bandwidth.




Figure 11: Output of the functional mean-shift algorithm on the 40 neural activity curves with the largest candidate bandwidth, h_2 . The plots are on the same scale. Each panel represents an individual cluster of neural activity curves and the bold line is the estimated modal curve for that cluster. With this choice of the bandwidth, the functional mean-shift algorithm detects that there are two bundles of curves that share a similar shape (top blue cluster) and a third cluster corresponding to curves with a more pronounced sigmoid shape (bottom yellow cluster).




Figure 12: Output of the functional mean-shift algorithm on the 40 neural activity curves with the smallest candidate bandwidth, h_1 . The plots are on the same scale. Each panel represents an individual cluster of neural activity curves and the bold line is the estimated modal curve for that cluster. From the plots, it appears that the choice of the bandwidth is just right: the functional mean-shift algorithm detects three internally homogeneous clusters. With this choice of the bandwidth, the algorithm identifies 2 atomic clusters (potential outliers) which correspond to the dashed curves in the two lower panels.

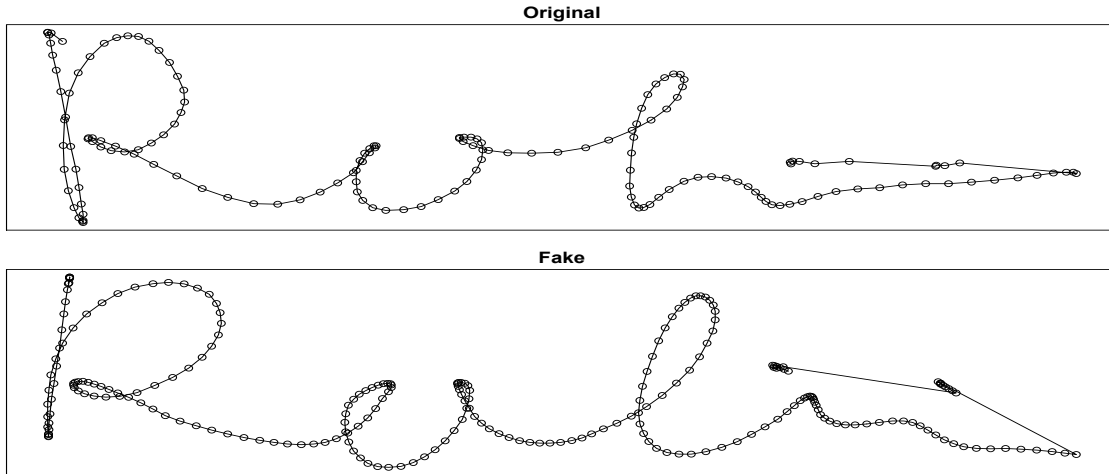


Figure 13: Examples of signatures from a sample of 40 SVC 2004 signatures (20 original, 20 fake). The upper panel displays an example of the owner's signature, while the bottom panel contains an example of a forged signature.

which captures the amount of acceleration along the line tangent to the signature X at time t .

From now on we identify the sample of signatures with the i.i.d. sample of tangential accelerations $\mathcal{S} = (S_1, S_2, \dots, S_n)$, $S_1 \sim P$. It is reasonable to imagine that if all the signatures in \mathcal{S} are produced by the same author (the owner of the signature), then the common distribution P of the S 's is 'unimodal'. In fact, in this case we might expect the distribution of the tangential accelerations to concentrate around a central tangential acceleration curve in the space of the S 's (the 'typical' tangential acceleration of the owner's signature); the scattering of the other tangential accelerations around this central curve would correspond to the random variation among different instances of an original signature. However, if some of the signatures in \mathcal{S} are instead forged by one or more skilled fakers whose forgeries appear to be good replicates of the original owner's signature to the naked eye, but whose tangential acceleration curves are sufficiently different from the tangential acceleration curves of the owner, then we may expect P to be 'multimodal'.

Suppose that we do not know exactly how many different authors produced the signatures in \mathcal{S} , i.e. we suspect that some signatures could be forged versions of the owner's signature. Running the functional mean-shift algorithm on \mathcal{S} would output of a set $\tilde{\mathcal{M}} = \{\tilde{\mu}_1, \dots, \tilde{\mu}_r\}$ of modal tangential accelerations (corresponding to local modes of \tilde{p} , the estimated surrogate density of P) and would assign each observed tangential acceleration $S \in \mathcal{S}$ to the closest mode in $\tilde{\mathcal{M}}$, thus allowing us to divide the sample signatures in r distinct clusters. If the functional mean-shift algorithm detects a unique modal signature, and therefore all the observed signatures are grouped in a unique cluster, we may conclude that there is no evidence (at least on the basis of the tangential accelerations) that there exist forged signatures in \mathcal{S} . However, if $\tilde{\mathcal{M}}$ contains more than a single modal curve ($r > 1$), and thus the functional mean-shift algorithm partitions \mathcal{S} in more than a single cluster, then we have empirical evidence that the sample \mathcal{S} may contain signatures both from the owner and (potentially) other authors. In this case, we may conclude that some signatures in \mathcal{S} are forged versions of the owner's original signature.

In this section, we test the functional mean-shift on a set of 40 signatures from a sample catalog of the SVC 2004 Signature Verification Competition (<http://www.cse.ust.hk/svc2004/download.html>, Sample Data, User 1). While it is known that in this sample there are 20 original owner's signatures (corresponding to the first 20 curves in the dataset) and 20 fake signatures (corresponding to the curves numbered from 21 to 40 in the dataset), we pretend that we do not know this information a priori. If the two groups of signatures are sufficiently separated in the space of the tangential accelerations, then the functional mean-shift should detect at least two modal curves. Our analysis represents an unsupervised counterpart of the analysis of [Geenens \(2011\)](#), who develops a nonparametric functional classifier for these data.

Figure 13 displays two signatures (one original, one fake) from the SVC 2004 sample. It is not immediately evident to the naked eye that the two signatures are not produced by the same author. Our first step is to smooth the x and y components of the 40 signatures in the sample with a locally quadratic polynomial to get a smooth representation of x' , x'' , y' and y'' . In particular, we use the `locpoly` function of the `KernSmooth` R package to perform the local polynomial smoothing and the function `dpill` of the same package to select the bandwidth parameter. With these smooth estimates of the derivatives, we then obtain




Figure 14: Smooth tangential accelerations obtained from the balanced dataset of 40 SVC 2004 signatures. The top panel displays the smooth tangential accelerations of all the 40 signatures together, whereas the middle and the bottom panels display the smooth tangential accelerations for the original and for the fake signatures, respectively.

the smooth tangential acceleration curve of equation (24) for each of the 40 signatures in the dataset which is then normalized so that it has unit L_2 norm. Figure 14 depicts the normalized smooth tangential accelerations obtained from the smooth representation of x' , x'' , y' and y'' and equation (24).

We apply the functional mean-shift algorithm with the asymmetric truncated Gaussian profile on the smooth tangential acceleration curves. Within the kernel, d is set to be the distance induced by the L_2 norm (as in [Geenens, 2011](#)), and the same heuristic of Section 8 is used to select two candidate bandwidths, $h_1 = 0.4485 \cdot \max_{X, X'} d(X, X')$ and $h_2 = 0.5242 \cdot \max_{X, X'} d(X, X')$. The plot of non-atomic clusters as a function of bandwidth for the tangential acceleration curves is depicted in Figure 15.

Figure 16 displays the output of the functional mean-shift on the smooth tangential acceleration curves when $h = h_1$ and Figure 17 shows the output of the algorithm when $h = h_2$. The functional mean-shift algorithm finds 5 clusters with $h = h_1$: the top cluster (black) contains 8 original signatures and appears very similar to the second, the fourth, and the fifth clusters (red, blue and yellow). These clusters contain 6 original and 1 fake signature (red), 2 original and 1 fake signature (blue), and 2 original signatures (yellow), respectively. The middle cluster (green) contains 15 of the 20 fake signatures. Overall, it seems that the clustering may be too fine. Figure 17 displays the output of the functional mean-shift algorithm when $h = h_2$. This time, the algorithm produces two internally homogeneous clusters which appear to meaningfully summarize the structure of the data: the top cluster (black) contains 19 of the 20 original signatures and the second cluster (green) contains all the 20 fake signatures.




Figure 15: Number of non-atomic clusters as a function of the bandwidth parameter, h . The number of non-atomic clusters stabilizes in the 2 highlighted ranges. The candidate bandwidth parameters are then set to be the midpoint of these 2 relevant ranges. The number in parentheses above the candidate bandwidths denote the number of sample curves that are clustered in non-atomic clusters using the corresponding candidate bandwidth.

10. Discussion

The mean-shift algorithm is an iterative algorithm that corresponds to a form of adaptive gradient ascent. The mean-shift algorithm has been extensively studied for scalar and vector data, but a counterpart of this algorithm for infinite-dimensional data has not been devised. In this paper, we introduce the *functional mean-shift algorithm*, a functional version of the mean-shift algorithm of [Fukunaga and Hostetler \(1975\)](#) and [Cheng \(1995\)](#) which is designed to deal with data that are intrinsically infinite-dimensional.

The functional mean-shift algorithm expands the applicability of the mean-shift algorithm well beyond the familiar Euclidean case. In [Section 5](#), we show that the functional mean-shift algorithm corresponds to a form of adaptive gradient ascent on the estimated surrogate density of random variables that are valued in an infinite-dimensional Hilbert space: for practitioners, this fact establishes a reassuring formal analogy with the standard mean-shift algorithm for scalar and vector data. In [Section 7](#), we provide a bootstrap-based simultaneous significance test for the local modes in infinite dimensions. We illustrate by means of a simulation that the test allows us to infer which of the candidate local modes identified by the functional mean-shift algorithm (and therefore which of the clusters) are significant and correspond to real features of the data generating process. In [Sections 8](#) and [9](#), we consider two examples of application which demonstrate the effectiveness of the algorithm for the task of clustering smooth curves.

A strength of the algorithm lies in the fact that it does not require the application of a dimension reduction procedure on the sample curves, nor does it require the user to




Figure 16: Output of the functional mean-shift algorithm on the 40 smooth tangential acceleration curves when $h = h_1$. The plots are on the same scale. Each panel represents an individual cluster of tangential acceleration curves and the bold line is the estimated modal curve for that cluster. The candidate local mode of the black cluster (top panel) contains 8 original signatures and appears very similar to that of the red (6 original, 1 fake), blue (2 original, 1 fake) and yellow (2 original) clusters (respectively, the second, the fourth and the fifth plots from the top). This similarity suggests that h_1 is producing a very fine clustering of the curves. The green cluster (mid plot) contains 15 of the 20 fake signatures. The remaining 5 atomic clusters are not shown in this figure.




Figure 17: Output of the functional mean-shift algorithm on the 40 smooth tangential acceleration curves when $h = h_2$. The plots are on the same scale. Each panel represents an individual cluster of tangential acceleration curves and the bold line is the estimated modal curve for that cluster. The black cluster (top) contains 19 of the 20 original signatures, whereas the green cluster (bottom) contains the 20 fake signatures. With this choice of the bandwidth, the functional mean-shift algorithm achieves a perfect separation between original and fake signatures. The algorithm finds an atomic cluster (dashed curve in the top panel, corresponding to the first original signature in the dataset), which is therefore flagged as a potential outlier.

choose the number of clusters a priori (in contrast to PCA-based clustering and k -means). Furthermore, the functional mean-shift algorithm is a flexible algorithm because it allows the user to tune the distance function and therefore to choose implicitly in which space to embed the data.

As it is the case with scalar or vector data, the effectiveness of the functional mean-shift algorithm comes at the cost of being able to properly select the bandwidth parameter h . The selection of the bandwidth is a crucial step in any analysis based on the mean-shift algorithm: a poor selection of the bandwidth can drive the algorithm to produce puzzling output. Several automatic and data-driven bandwidth selectors have been proposed for the mean-shift algorithm in the Euclidean setting, some of which reduce to the problem of optimally choosing the bandwidth parameter for the estimation of the underlying probability density (see [Park and Marron, 1990](#) and [Sheather and Jones, 1991](#) who develop optimal *plug-in* bandwidth selectors or, among others, [Faraway and Jhun, 1990](#) for a bootstrap-based bandwidth selector). This is arguably not the optimal strategy to select the bandwidth for the mean-shift algorithm: the mean-shift algorithm is a gradient ascent algorithm, therefore it appears more natural to select the bandwidth in such a way to optimize the estimation of the density derivatives ([Chacón and Duong, 2013](#)). An incomplete list of other proposals includes [Comaniciu et al. \(2001\)](#), who propose strategies to select data-driven adaptive bandwidths, [Einbeck \(2011\)](#), who introduces a bandwidth selector based on *principal curves* ([Hastie and Stuetzle, 1989](#); [Flury, 1990](#); [Flury, 1993](#)), and [Genovese et al. \(2013\)](#), who propose a bandwidth selection method based on the maximization of the number of significant modes.

The significance test of Section 7 allows us to choose the bandwidth in such a way to maximize the number of significant modes as suggested in [Genovese et al. \(2013\)](#). Also, the multi-bandwidth analysis that we discuss in Sections 8 and 9 offers an alternative approach for the selection of the bandwidth in applied work. However, it seems harder to come up with a theory for bandwidth selection in the infinite-dimensional setting that we consider. We regard the investigation of theoretically justified data-driven bandwidth selectors for the functional mean-shift algorithm as an interesting direction for further theoretical research.

Acknowledgments

ADD ACKNOWLEDGMENTS

Proof of Lemmata

Proof of Lemma 4

The Gateaux differential of the estimated surrogate density can be obtained by applying the chain rule. We report here a more explicit calculation that highlights the relationship between the kernel K used in the functional mean-shift algorithm and its shadow G that is used in the corresponding gradient ascent analog.

$$\begin{aligned}
 \tilde{p}_x(y) &= \left. \frac{d}{d\alpha} \tilde{p}(x + \alpha y) \right|_{\alpha=0} = \left. \frac{d}{d\alpha} w_G(\mathcal{S}) \sum_{X \in \mathcal{S}} G_h(X, x + \alpha y) \right|_{\alpha=0} \\
 &= \left. \frac{d}{d\alpha} w_G(\mathcal{S}) \sum_{X \in \mathcal{S}} g\left(\frac{d(X, x + \alpha y)}{h(X)}\right) \right|_{\alpha=0} \\
 &= \left. \frac{d}{d\alpha} w_G(\mathcal{S}) \sum_{X \in \mathcal{S}} g\left(\frac{\|X - x - \alpha y\|}{h(X)}\right) \right|_{\alpha=0} \\
 &= w_G(\mathcal{S}) \sum_{X \in \mathcal{S}} -\frac{1}{h(X)} g'\left(\frac{\|X - x - \alpha y\|}{h(X)}\right) \frac{\alpha \|y\|^2 - \langle X - x, y \rangle}{\|X - x - \alpha y\|} \Big|_{\alpha=0} \\
 &= \left\langle y, w_G(\mathcal{S}) \sum_{X \in \mathcal{S}} -\frac{1}{h(X)} g'\left(\frac{d(X, x)}{h(X)}\right) \frac{X - x}{d(X, x)} \right\rangle \\
 &= \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} k\left(\frac{d(X, x)}{h(X)}\right) (X - x) \right\rangle \\
 &= \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x) (X - x) \right\rangle.
 \end{aligned}$$

Proof of Lemma 11

$$\begin{aligned}
 \tilde{p}_x^{(2)}(y, z) &= \frac{d}{d\alpha} \tilde{p}_{x+\alpha z}(y) \Big|_{\alpha=0} \\
 &= \frac{d}{d\alpha} \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} K_h(X, x + \alpha z) (X - x - \alpha z) \right\rangle \Big|_{\alpha=0} \\
 &= \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \frac{d}{d\alpha} K_h(X, x + \alpha z) (X - x - \alpha z) \right\rangle \Big|_{\alpha=0} \\
 &= \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[-K_h(X, x + \alpha z) z \right. \right. \\
 &\quad \left. \left. - \frac{1}{h(X)} \frac{\langle X - x, z \rangle + \alpha \|z\|^2}{\|X - x - \alpha z\|} K'_h(X, x + \alpha z) (X - x - \alpha z) \right] \right\rangle \Big|_{\alpha=0} \\
 &= \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[-\frac{1}{h(X)} \frac{\langle X - x, z \rangle}{\|X - x\|} K'_h(X, x) (X - x) \right. \right. \\
 &\quad \left. \left. - K_h(X, x) z \right] \right\rangle \\
 &= - \left\langle y, Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} \frac{\langle X - x, z \rangle}{\|X - x\|} K'_h(X, x) (X - x) \right. \right. \\
 &\quad \left. \left. + K_h(X, x) z \right] \right\rangle \\
 &= -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) \frac{\langle X - x, y \rangle \langle X - x, z \rangle}{\|X - x\|} \right. \\
 &\quad \left. + K_h(X, x) \langle y, z \rangle \right].
 \end{aligned}$$

Proof of Lemma 12

$$\begin{aligned}
 \tilde{\lambda}_x &= \sup_{\|y\|=1} \tilde{p}_x^{(2)}(y, y) = \\
 &= \sup_{\|y\|=1} -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) \frac{\langle X - x, y \rangle \langle X - x, y \rangle}{\|X - x\|} \right. \\
 &\quad \left. + K_h(X, x) \langle y, y \rangle \right] \\
 &= \sup_{\|y\|=1} -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) \frac{\|X - x\|^2 + \|y\|^2 + 2\langle X - x, y \rangle}{\|X - x\|} \right. \\
 &\quad \left. + K_h(X, x) \|y\|^2 \right] \\
 &= -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) [\|X - x\| + \|X - x\|^{-1}] + K_h(X, x) \right] \\
 &\quad + \sup_{\|y\|=1} -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{2}{h^3(X)} K'_h(X, x) \frac{\langle X - x, y \rangle}{\|X - x\|} \\
 &= -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) [\|X - x\| + \|X - x\|^{-1}] \right. \\
 &\quad \left. + K_h(X, x) \right] + \sup_{\|y\|=1} \left\langle y, -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{2}{h^3(X)} K'_h(X, x) \frac{X - x}{\|X - x\|} \right\rangle \\
 &= -Cw_G(\mathcal{S}) \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) [\|X - x\| + \|X - x\|^{-1}] \right. \\
 &\quad \left. + K_h(X, x) \right] + 2Cw_G(\mathcal{S}) \left\| \sum_{X \in \mathcal{S}} \frac{1}{h^3(X)} K'_h(X, x) \frac{X - x}{\|X - x\|} \right\| \\
 &= Cw_G(\mathcal{S}) \left[2 \left\| \sum_{X \in \mathcal{S}} \frac{1}{h^3(X)} K'_h(X, x) \frac{X - x}{\|X - x\|} \right\| - \right. \\
 &\quad \left. \sum_{X \in \mathcal{S}} \frac{1}{h^2(X)} \left[\frac{1}{h(X)} K'_h(X, x) (\|X - x\| + \|X - x\|^{-1}) + K_h(X, x) \right] \right].
 \end{aligned}$$

References

- Youness Aliyari Ghassabeh. On the convergence of the mean shift algorithm in the one-dimensional space. *Pattern Recognition Letters*, 34(12):1423 – 1427, 2013.
- Antonio Ambrosetti and Giovanni Prodi. *A primer of nonlinear analysis*. Number 34. Cambridge University Press, 1995.
- Ery Arias-Castro, David Mason, and Bruno Pelletier. On the estimation of the gradient lines of a density and the consistency of the mean-shift algorithm. *unpublished manuscript*, 2013.
- Adelchi Azzalini and Nicola Torelli. Clustering via nonparametric density estimation. *Statistics and Computing*, 17(1):71–80, 2007.
- Denis Bosq. *Linear processes in function spaces: theory and applications*. Springer, 2000.
- Miguel Á. Carreira-Perpiñán. Fast nonparametric clustering with gaussian blurring mean-shift. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 153–160, 2006.
- José E. Chacón. Clusters and water flows: a novel approach to modal clustering through morse theory. *arXiv preprint arXiv:1212.1384*, 2012.
- José E. Chacón. A population background for nonparametric density-based clustering. *unpublished manuscript*, 2014.
- José E. Chacón and Tarn Duong. Data-driven density derivative estimation, with applications to nonparametric clustering and bump hunting. *Electronic Journal of Statistics*, 7: 499–532, 2013.
- Frédéric Chazal, Leonidas J Guibas, Steve Y Oudot, and Primoz Skraba. Persistence-based clustering in riemannian manifolds. *Journal of the ACM (JACM)*, 60(6):41, 2013.
- Yizong Cheng. Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):790–799, 1995.
- Herman Chernoff. Estimation of the mode. *Annals of the Institute of Statistical Mathematics*, 16(1):31–41, 1964.
- Dorin Comaniciu and Peter Meer. Mean shift: a robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24:603–619, 2002.
- Dorin Comaniciu, Visvanathan Ramesh, and Peter Meer. The variable bandwidth mean shift and data-driven scale selection. In *Proceedings of the Eighth IEEE International Conference on Computer Vision*, volume 1, pages 438–445, 2001.
- Antonio Cuevas and Ricardo Fraiman. A plug-in approach to support estimation. *The Annals of Statistics*, 25(6):2300–2312, 1997.

- Sophie Dabo-Niang, Frédéric Ferraty, and Philippe Vieu. Estimation du mode dans un espace vectoriel semi-normé. *Comptes Rendus Mathématique*, 339(9):659–662, 2004.
- Aurore Delaigle and Peter Hall. Defining probability density for a distribution of random functions. *The Annals of Statistics*, 38(2):1171–1193, 2010.
- David L. Donoho and Richard C. Liu. Geometrizing rates of convergence, iii. *The Annals of Statistics*, pages 668–701, 1991.
- William F. Eddy. Optimum kernel estimators of the mode. *The Annals of Statistics*, 8(4):870–882, 1980.
- Jochen Einbeck. Bandwidth selection for mean-shift based unsupervised learning techniques: a unified approach via self-coverage. *Journal of Pattern Recognition Research*, 6(2):175–192, 2011.
- Julian J. Faraway and Myoungshic Jhun. Bootstrap choice of bandwidth for density estimation. *Journal of the American Statistical Association*, 85(412):1119–1122, 1990.
- Frédéric Ferraty and Philippe Vieu. *Nonparametric functional data analysis: theory and practice*. Springer, 2006.
- Frédéric Ferraty, Ali Laksaci, and Philippe Vieu. Estimating some characteristics of the conditional distribution in nonparametric functional models. *Statistical Inference for Stochastic Processes*, 9(1):47–76, 2006.
- Frédéric Ferraty, Nadia Kudraszow, and Philippe Vieu. Nonparametric estimation of a surrogate density function in infinite-dimensional spaces. *Journal of Nonparametric Statistics*, 24(2):447–464, 2012.
- Bernard D. Flury. Estimation of principal points. *Applied Statistics*, pages 139–151, 1993.
- Bernhard A. Flury. Principal points. *Biometrika*, 77(1):33–41, 1990.
- Keinosuke Fukunaga and Larry Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on Information Theory*, 21(1):32–40, 1975.
- Theo Gasser, Peter Hall, and Brett Presnell. Nonparametric estimation of the mode of a distribution of random curves. *Journal of the Royal Statistical Society, Series B*, 60(4):681–691, 1998.
- Gery Geenens. A nonparametric functional method for signature recognition. In Frédéric Ferraty, editor, *Recent advances in functional data analysis and related topics*, Contributions to statistics, pages 141–147. Physica-Verlag HD, 2011.
- Christopher Genovese, Marco Perone-Pacífico, Isabella Verdinelli, and Larry Wasserman. Nonparametric inference for density modes. *arXiv preprint arXiv:1312.7567*, 2013.
- John A. Hartigan. *Clustering algorithms*. John Wiley & Sons, 1975.

- Trevor Hastie and Werner Stuetzle. Principal curves. *Journal of the American Statistical Association*, 84(406):502–516, 1989.
- Lajos Horváth and Piotr Kokoszka. *Inference for functional data with applications*. Springer, 2012.
- Julien Jacques and Cristian Preda. Functional data clustering: a survey. *Advances in Data Analysis and Classification*, pages 1–25, 2013.
- Jussi Klemelä. Adaptive estimation of the mode of a multivariate density. *Journal of Nonparametric Statistics*, 17(1):83–105, 2005.
- Jing Lei, Alessandro Rinaldo, and Larry Wasserman. A conformal prediction approach to explore functional data. *Annals of Mathematics and Artificial Intelligence*, pages 1–15, 2013.
- Jia Li, Surajit Ray, and Bruce G. Lindsay. A nonparametric statistical approach to clustering via mode identification. *Journal of Machine Learning Research*, 8(8):1687–1723, 2007a.
- Xiangru Li, Zhanyi Hu, and Fuchao Wu. A note on the convergence of the mean shift. *Pattern Recognition*, 40(6):1756 – 1762, 2007b.
- Byeong U. Park and James S. Marron. Comparison of data-driven bandwidth selectors. *Journal of the American Statistical Association*, 85(409):66–72, 1990.
- Emanuel Parzen. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3):1065–1076, 1962.
- James O. Ramsay and Bernard W. Silverman. *Functional Data Analysis*. Springer, 2005.
- Alessandro Rinaldo and Larry Wasserman. Generalized density clustering. *The Annals of Statistics*, 38(5):2678–2722, 2010.
- Alessandro Rinaldo, Aarti Singh, Rebecca Nugent, and Larry Wasserman. Stability of density-based clustering. *Journal of Machine Learning Research*, 13(1):905–948, 2012.
- Joseph P. Romano. On weak convergence and optimality of kernel density estimates of the mode. *The Annals of Statistics*, 16(2):629–647, 1988.
- Simon J. Sheather and Michael C. Jones. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society, Series B*, 53(3): 683–690, 1991.
- Bernard W. Silverman. Using kernel density estimates to investigate multimodality. *Journal of the Royal Statistical Society, Series B*, pages 97–99, 1981.
- Bernard W. Silverman. *Density estimation for statistics and data analysis*. Chapman & Hall, 1986.

- Werner Stuetzle. Estimating the cluster tree of a density by analyzing the minimal spanning tree of a sample. *Journal of Classification*, 20(1):025–047, 2003.
- Werner Stuetzle and Rebecca Nugent. A generalized single linkage method for estimating the cluster tree of a density. *Journal of Computational and Graphical Statistics*, 19(2), 2010.
- Philippe Vieu. A note on density mode estimation. *Statistics & Probability Letters*, 26(4): 297–307, 1996.
- Vladimir Vovk, Ilia Nouretdinov, and Alex Gammerman. On-line predictive linear regression. *The Annals of Statistics*, 37(3):1566–1590, 2009.